Network Monitoring today: why, how, challenges, infrastructures, federations and the Grid

Prepared by Les Cottrell, SLAC, for the

Grid Performance Workshop

UCL, May 12-13, 2004

www.slac.stanford.edu/grp/scs/net/talk03/gridperf-may04.ppt

Partially funded by DOE/MICS Field Work Proposal on Internet End-to-end Performance Monitoring (IEPM), also supported by IUPAP

Why
(Can't manage what you can't measure)

- Need measurements for both production networks & tesbeds:
  - Planning, setting expectations, policy/funding
  - Trouble-shooting: reliability & performance
    - Problems may not be logical, e.g. most Internet problems caused by operator error (Sci Am Jun'03), most LAN problems are Ethernet duplex, host config, bugs
    - Made hard by transparency, size & rate of change of network
    - *A distributed system is one in which I can't get my work done because a computer I never heard of has failed.* Butler Lampson
  - Application steering (e.g. Grid data replication)
- E2E performance problem is THE critical user metric

E.g. Policy - trends

S.E. Europe, Russia: **catching up**
Latin Am., Mid East, China: **keeping up**
India, Africa: **falling behind**

C. Asia, Russia, S.E. Europe, L. America, M. East, China: 4-5 yrs behind

India, Africa: 7 yrs behind

Important for policy makers

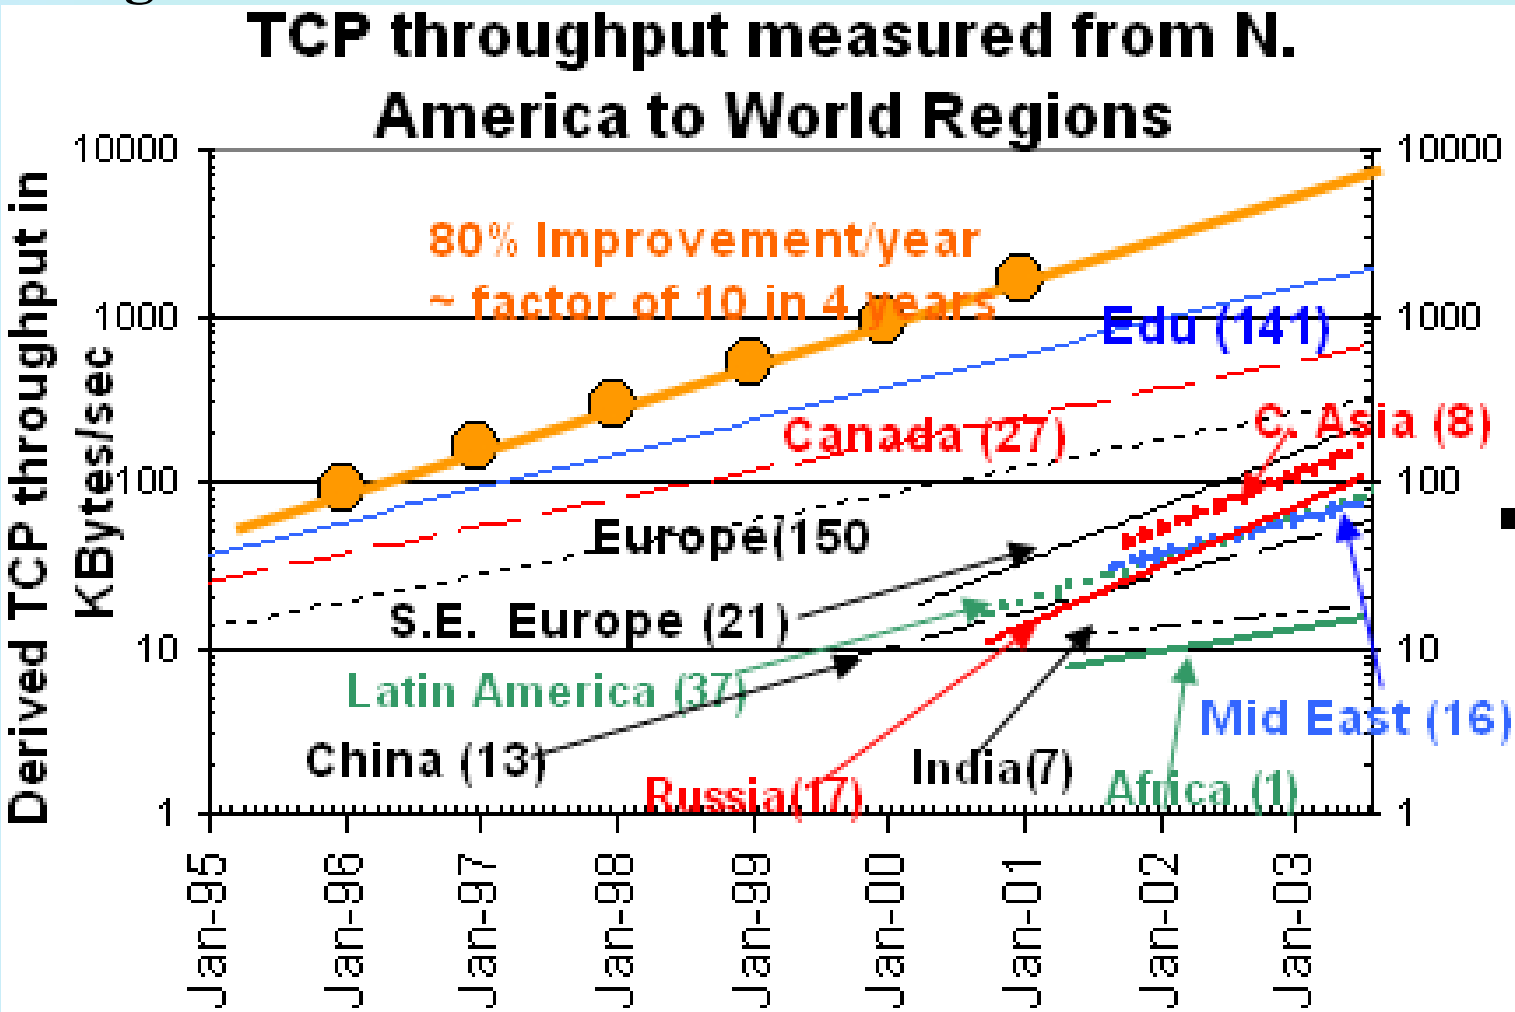

E.g. Changes in network topology (BGP) result in dramatic change in performance

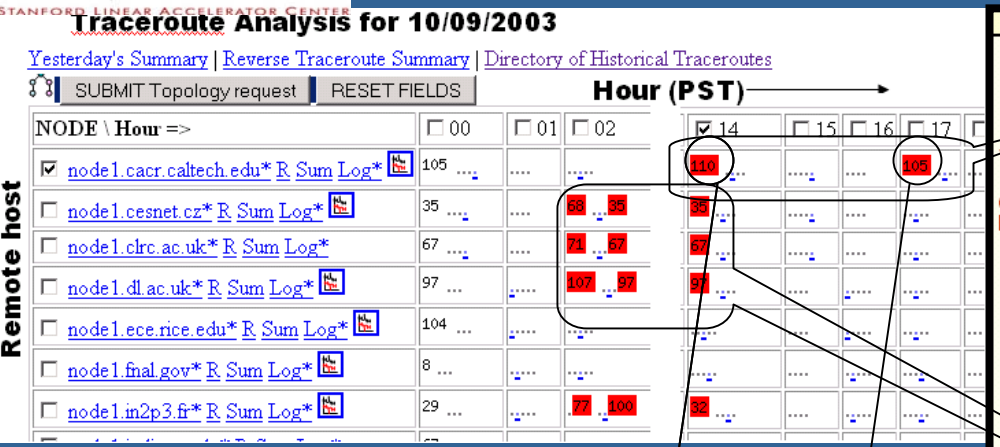

Snapshot of traceroute summary table

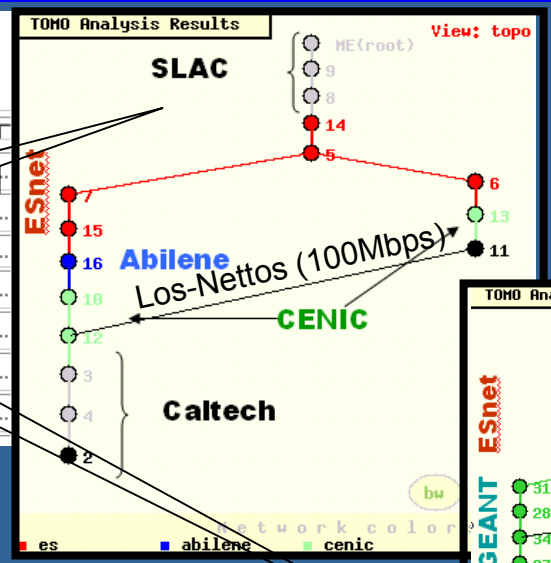

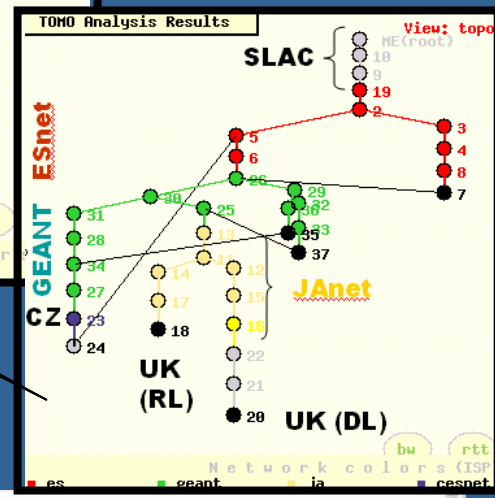

Samples of traceroute trees generated from the table

Notes:

1. Caltech misrouted via Los-Nettos 100Mbps commercial net 14:00-17:00
2. ESnet/GEANT working on routes from 2:00 to 14:00
3. A previous occurrence went un-noticed for 2 months
4. Next step is to auto detect and notify

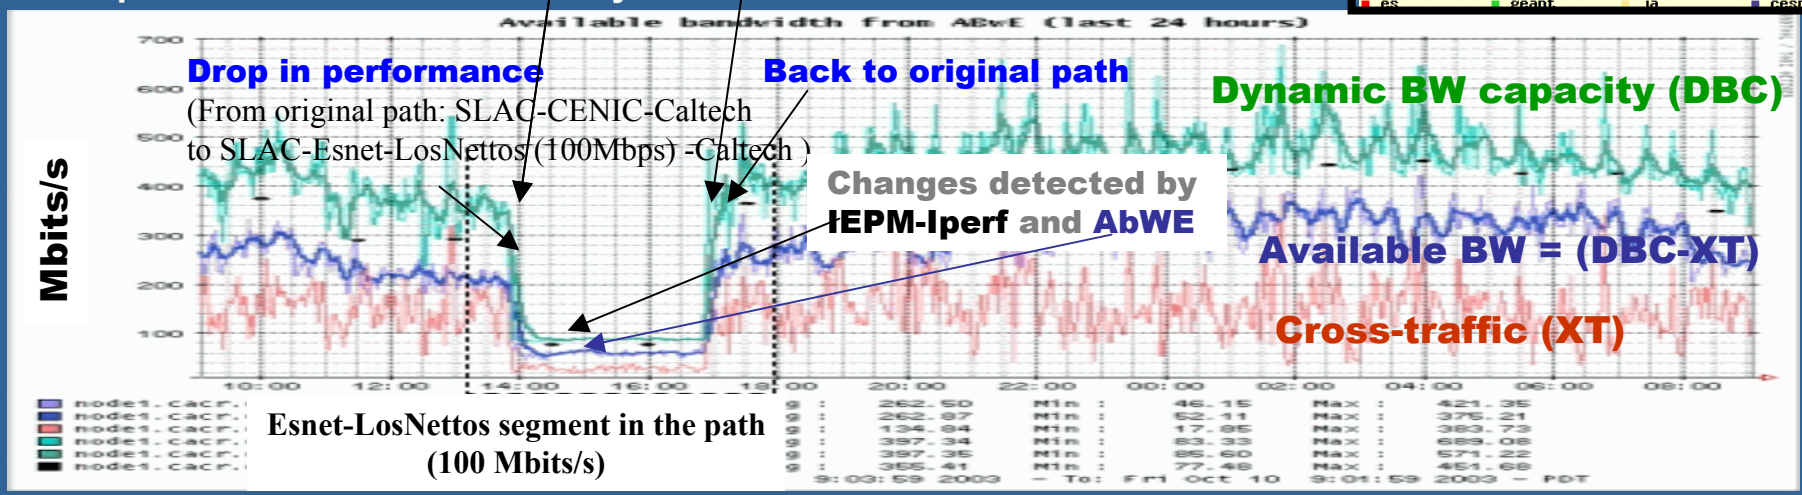

ABwE measurement one/minute for 24 hours Thurs Oct 9 9:00am to Fri Oct 10 9:01am

Methods

- Active Measurement probes:
  - Include: Ping, traceroute, owamp, pathload/abwe, major apps (e.g. bbftp, bbcp, GridFTP…)
  - Typically used for end-to-end testing
  - Injects data into network, can be non-negligible
- Passive tools:
  - Include: SNMP, NetFlow, OCxMon, NetraMet, cflowd, SCNM
  - Typically used at border or inside backbones
    - SNMP heavily used for utilization, errors on LAN & backbones
    - Flows for traffic characterization and intrusion detection
  - Need access to network devices (e.g. routers, taps)
    - Can generate a lot of data
- Need to put together data from multiple sources
  - Different probes, different source & destinations, network-centric & end-to-end

Some Challenges for Active monitoring

- Bandwidth used, e.g. iperf etc. & apps
  - Sampling rate (Nyquist's theorem),
  - Relevance to application needs
  - Measure loss to 10% on a path with 1 in 10K loss requires a million pings
- For TCP tools: configuring windows at clients/servers and optimizing windows, streams
- Some lightweight tools (e.g. packet pairs) not effective at >> 1Gbits/s
- Many tools tuned for shared TCP/IP nets not for dedicated circuits
- Simplifying use and understanding for end-user
- Automating problem detection & resolution,

Network Impact

- Heavyweight: iperf, bbcp, bbftp, GridFTP (IEPM-BW, PiPES …)…
  - Noticeable impact, run infrequently (e.g. hourly) , and for short time (e.g. tens of seconds), only small number of sites
  - Need scheduling
  - Close to what applications see
- Lightweight: Ping, traceroute, ABwE etc.
  - E.g. PingER, AMP
  - Can do on demand, no need to set things up in advance (no server to install), no scheduling needed, can involve thousands of sites
- Medium weight (ABwE, pathload etc.)
  - E.g. IEPM-LITE, Scriptroute
  - Needs server/mirror install, low traffic (ABwE 1kbps avg), no scheduling

Infrastructures

- Many measurement projects with different emphases, different communities
  - Passive (usually requires network control, used at borders and on backbones, e.g. MICSmon/Netflow, ISP/SNMP, SCNM)
  - Active: amount of network "pollution"
    - Lightweight (PingER, AMP, Surveyor, RIPE …)
    - Medium weight (PiPES, NWS, IEPM-Lite …)
    - Heavy weight/hi-perf (IEPM-BW, NTAF
  - End-to-end vs net centric (skitter, macroscopic views)
  - Repetitive (PingER, AMP, IEPM, PiPES, NWS, NTAF, …)
  - On demand, or non-production (NDT, NIMI, PiPES …)
  - Dedicated hardware (AMP, RIPE, NDT, PlanetLab …)
  - Hierarchical (e.g. AMP) vs Full mesh (e.g. PingER)
- For a table comparing 13 public domain infrastructures, see: www.slac.stanford.edu/grp/scs/net/proposals/infra-mon.html

NMI challenges

- Sustaining deployment/operation in multi-agency / international world
- Scaling beyond hundreds of hosts very hard over the long term:
  - Hosts change, upgrade, new OS
    - No control over shared hosts
      - Depend on friendly admin contacts who may be busy, uninterested, have moved etc.
    - Policy/fears at remote site can make dedicated changes painful
    - web100 upgrades not coordinated with Linux upgrades
    - New TCP kernel upgrades not coordinated with OS upgrades
  - Hosts age, become measurement bottleneck
    - Need constant upgrades for dedicated hosts
  - Probes (iperf etc.) change: new features, patches
  - Scheduling to prevent interference for heavyweight tests
- Appropriate security: keeping track of credentials, upgrade/patches, multiple-policies, port blocking

So Recognize

- Unrealistic to think multiple admin domains will all deploy one and the same infrastructure
  - Scaling and interests make unrealistic
- Multiple-domain, multi-infrastructures will be deployed
- Need to tie together heterogeneous collection of monitoring systems
  - Create a federation of existing NMIs
  - Infrastructures work together
  - Share data with peer infrastructures and others using a common set of protocols for describing, exchanging & locating monitoring data (e.g. GGF NMWG)
  - Enables much improved overall view of network using multiple measurement types from multiple sources

MAGGIE Proposal

- Measurement and Analysis for the Global Grid and Internet End-to-end performance
- Contribute to, utilize the GGF NMWG naming hierarchy and the schema definitions for network measurements
- Develop tools to allow sharing
  - Web services based
  - Integrate information from multiple sources
- Brings together several major infrastructure participants: LBNL (NTAP, SCNM), SLAC (IEPM-PingER/BW), Internet2 (PiPES, NDT), NCSC (NIMI), U Delaware, ESnet
- Will work with others, e.g. MonALISA, AMP, UltraLight, PPDG, StarLIght, UltraScienceNet

Federation goals

- Appropriate security
- Interoperable
- Useful for applications, network engineers, scientists & end users
- Easy to deploy & configure
- As un-intrusive as possible
- As accurate & timely as possible
- Identify most useful features of each NMI to improve each NMI faster than working alone

From measurements to the Grid

- Given measurements or the ability to make them, how is that useful to the Grid?
- Grid application needs to place or retrieve data with high performance and robustness
  - Maybe use multiple sites in parallel
    - Some similarities with P2P such as BitTorrent, eDonkey, Kazaa, Gnutella etc. chunking of files
    - But different goals
      - Grid few well-known sites known in advance, high-perf links, does not face legal troubles, free-riding etc. of P2P
- Need to find optimal site(s) to get data from based on expected achievable throughput
  - Can use existing measurements & predictions
  - Can make measurements on demand

Use Existing Measurements

- Need a way to discover “relevant” measurements
  - Between possible pairs or “closely” related pairs
- Need a request protocol/schema
- Need a response schema for results
- GGF NMWG are working on these issues.

On-demand

- Application somehow knows where chunks of data may be found
- Makes measurements of bandwidth from application site to chunk locations
  - Assumes have appropriate servers at chunk locations (e.g. ABwE reflector), or use ubiquitous server (e.g. ping)
- Uses this to choose initial locations to get a complete set of chunks from

Challenges

- Optimal chunk locations may change during transfer (chunk location may become inaccessible, or its performance may drop)
  - So need measurements during transfer
  - This may make it attractive to instrument the application so it can make its own measurements on the data being transferred
    - Need library to simplify modifying each application
- Throughput advantages of multiple parallel site transfers may be no better than multiple parallel streams between a well connected source & destination (may share the same bottleneck)
- Do network measurements relate to file transfer rates?

NMI Challenges:

- Reduce “Wizard gap”
- Applications cross agency AND international funding boundaries (includes Digital Divide)
- Incent multi-disciplinary teams, including people close to scientists, operational teams
  - Make sure what is produced is used, tested in real environment, include deployment in proposals
- *Network management research historically underfunded, because it is difficult to get funding bodies to recognize as legitimate networking research*, IAB
- Without excellent trouble-shooting capabilities, the Grid vision will fail

More Information

- Some Measurement Infrastructures:
  - CAIDA list: www.caida.org/analysis/performance/measinfra/
  - AMP: amp.nlanr.net/, PMA http://pma..nlanr.net
  - IEPM/PingER home site: www-iepm.slac.stanford.edu/
  - IEPM-BW site: www-iepm.slac.stanford.edu/bw
  - NIMI: ncne.nlanr.net/nimi/
  - RIPE: www.ripe.net/test-traffic/
  - NWS: nws.cs.ucsb.edu/
  - Internet2 PiPES: e2epi.internet2.edu/
- Tools
  - CAIDA measurement taxonomy: www.caida.org/tools/
  - SLAC Network Tools: www.slac.stanford.edu/xorg/nmtf/nmtf-tools.html
- Internet research needs:
  - www.ietf.org/internet-drafts/draft-iab-research-funding-00.txt
  - www.slac.stanford.edu/grp/scs/net/talk03/lsn-jun03.ppt

Automatic Step change Detection

- Too many graphs to review each morning!
- Motivated by drop in bandwidth between SLAC &Caltech
  - Started late August 2003
  - Reduced achievable throughput by factor of 5
  - Not noticed until October 2003
  - Caused by faulty routing over commercial network
  - After notifying ISP, it was fixed in 4 hours!
  - See http://www.slac.stanford.edu/grp/scs/net/case/caltech/ for details

SLAC Caltech achievable throughput April – November 2003

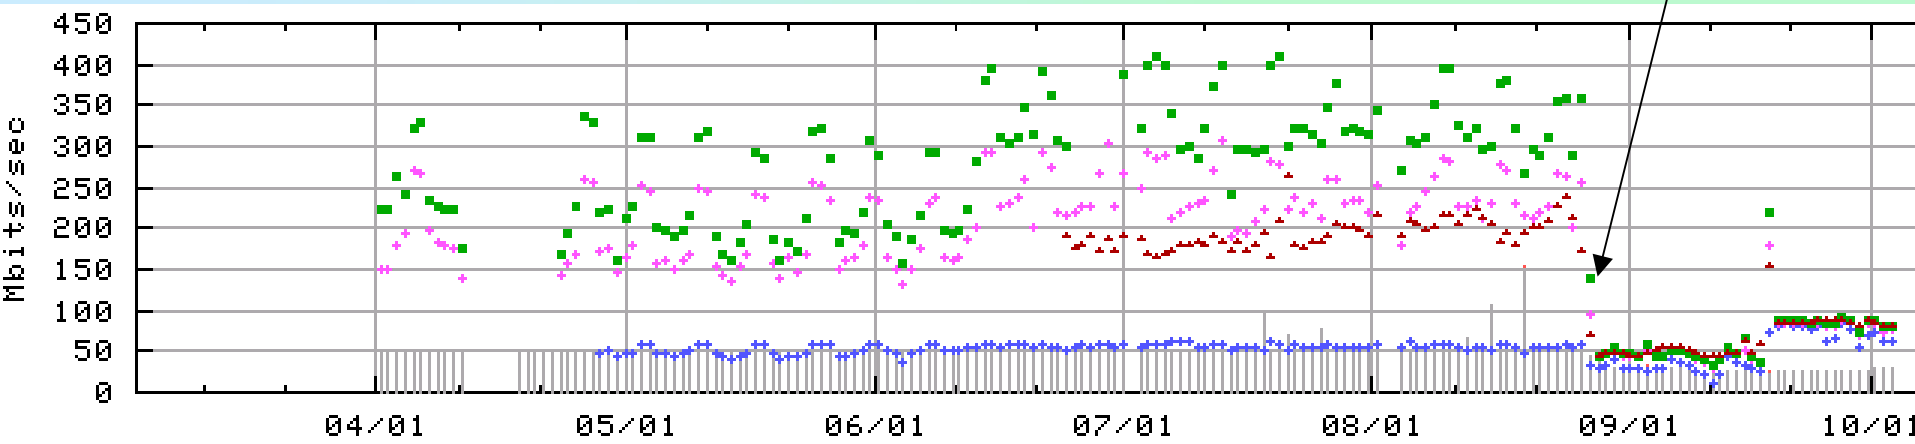

Automatic available bandwidth step change detection

- Still developing, evolving from earlier work:
  - Arithmetic weighted moving averages
  - NLANR “Plateau” algorithm work, see http://byerley.cs.waikato.ac.nz/~tonym/papers/event.pdf
- Goals catches important changes, with few false alerts

Plateau algorithm

- Roughly speaking:
  - Has a history buffer to describe past behavior
    - **History buffer duration** currently 600 mins
  - Plus a trigger buffer of data suggesting a change
    - **Trigger buffer duration** (evaluating typically 10-60 mins) indicates how long the change has to occur for
  - History mean (μ) and std. dev. (σ) use by trigger selector
    - If new_value outside μ +- **sensitivity**$*\sigma$ add to trigger buffer
    - If new_value outside μ +- 2***sensitivity**$*\sigma$ then also an outlier (don't add to stats)
    - Else goes in history buffer
  - Look for big difference in trigger and history buffer means
- Also looking at Principal Component Analysis (Crovella et al) of multi variables (e.g. capacity, cross-traffic, RTT …) which may also help with diurnal changes.

Examples

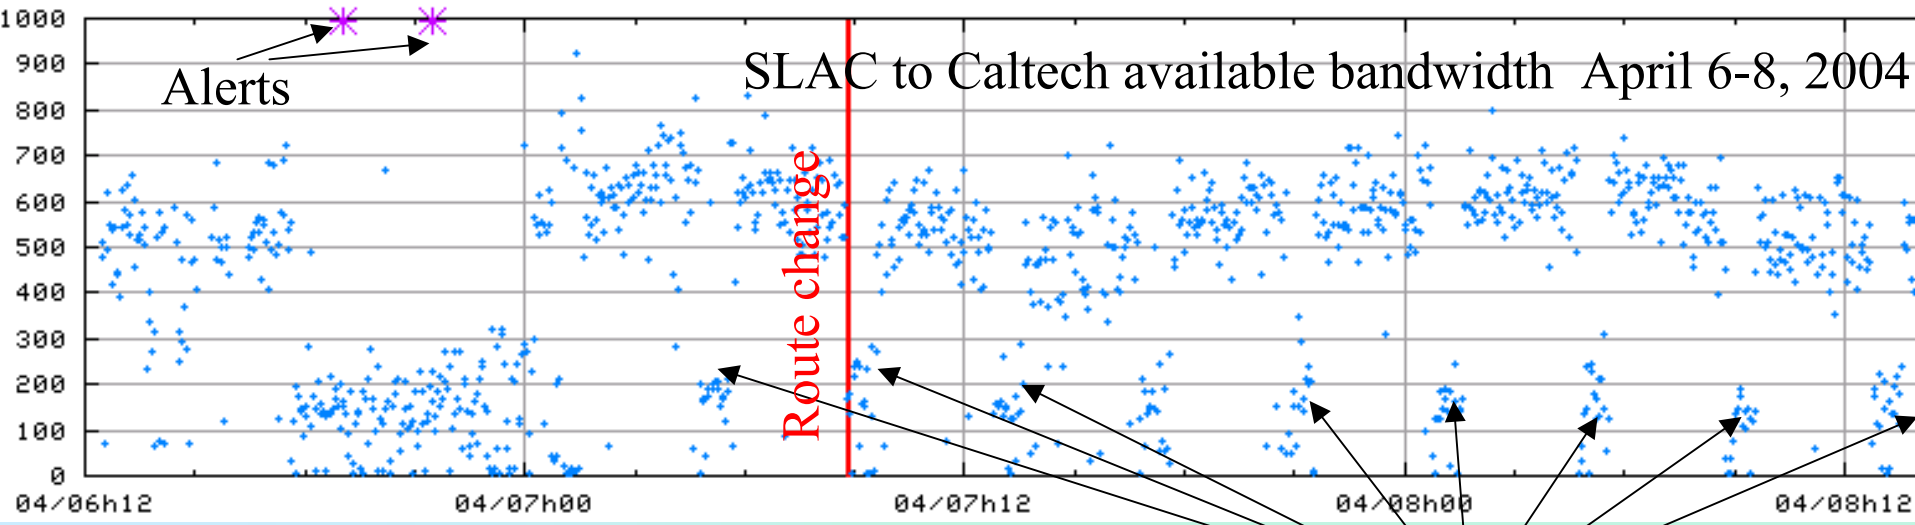

History duration: 600 mins, trigger duration: 30 mins, threshold: 40%, sensitivity: 2
With trigger duration: 60 only see one alert, with trigger duration: 10 catch alerts

Unreachable

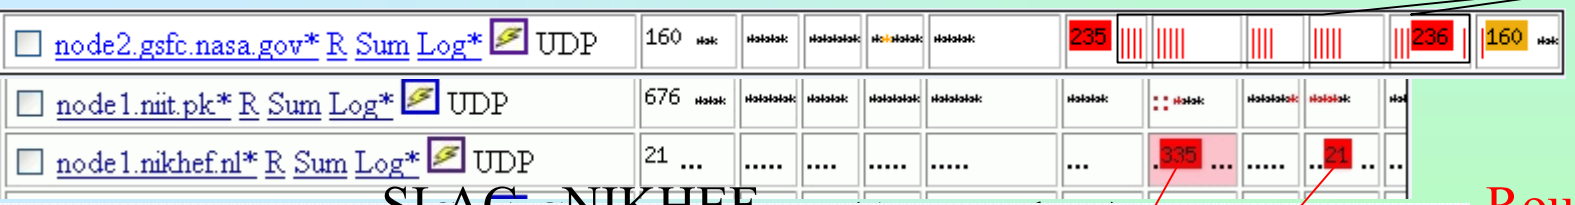

SLAC - NIKHEF

Route changes